NASA-CR-197040

NASA
7N-60-CR
30154
9P

FINAL REPORT FOR NCC2-408

RIACS SPARSE DISTRIBUTED MEMORY RESEARCH PROJECT

10/15/85 - 12/31/88

Principal Investigator: Pentti Kanerva

One of NASA's grand challenges is to build autonomous machines and systems that are capable of learning to perform tasks too tedious, or in places too remote and too hostile, for humans. The goals of the Learning Systems Division of RIACS is to find new approaches to autonomous systems based upon sound mathematical and engineering principles, to understand how information processing is organized in animals, and to test the applicability of these new approaches to the grand challenge. The research program includes the development of theory, implementations in software and hardware, and explorations of potential areas for applications. The Learning Systesm Division consists of two projects that pursue these goals, the Sparse Distributed Memory project and the Bayesian Learning project.

Now beginning its fourth year, the Sparse Distributed Memory (SDM) project is investigating the theory and applications of a massively parallel computing architecture that will support the storage and retrieval of sensory and motor patterns characteristic of autonomous systems. The project is centered in studies of the memory itself and in the use of the memory to solve problems in speech, vision, and robotics. Investigation of methods for encoding sensory data is an important part of the research. Examples of NASA missions that may benefit from this work are Space Station, planetary rovers, and solar exploration. Sparse distributed memory offers promising technology for systems that must learn through experience and that must be capable of adapting to new circumstances, and for operating any large complex system requiring automatic monitoring and control. This work, which is conducted primarily within RIACS, has included collaborations with NASA, Hewlett-Packard Corporation, MCC, and Stanford University.

Sparse distributed memory is a massively parallel architecture motivated by efforts to understand how the human brain works, given that the brain comprises billions of sparsely interconnected neurons, and by the desire to build machines capable of similar behavior. Sparse distributed memory is an associative memory, able to retrieve information from cues that only partially match patterns stored in the memory. It is able to store long temporal sequences derived from the behavior of a complex system, such as progressive records of the system's sensory data and corresponding records of the its motor controls. Using its records of successful behavior in the past, sparse distributed memory can be used to recognize a similar circumstance in the present and to *predict* appropriate responses. Unlike numerical and symbolic computers, sparse distributed memory is a pattern computer, designed to process very large patterns formulated as bit strings that may be thousands of bits long. Each such bit string can serve as both content and address within the memory. Our project is concerned with research into aspects of sparse distributed memory that will enable us to evaluate and build autonomous systems based upon sparse distributed memory.

In early 1988, Kanerva finished his book, *Sparse Distributed Memory*, and MIT Press's Bradford Books published it in late 1988. This book contains the seminal ideas of the memory and a chapter on the architecture of autonomous systems that learn.

(NASA-CR-197040) RIACS SPARSE DISTRIBUTED MEMORY RESEARCH PROJECT Final Report, 15 Oct. 1985 - 31 Dec. 1988 (Research Inst. for Advanced Computer Science) 9 p

N95-70491

Unclas

Z9/60 0030154

SYSTEM THEORY

Pentti Kanerva

An SDM-based model of an autonomous learning system has been reported, incorporating sensors, motors, memory, and a built-in preference function. The system has a central place called the *focus* that accounts for the system's subjective experience. The entity in the focus is a very large pattern that encodes everything about that moment--that is, any specific things that the system is attending to, the system's action, and the overall context. The memory is addressed by the focus, the memory's output goes into the focus, the senses feed into the focus, and the motors are driven from the focus. A system with such an architecture should be capable of learning how the world works and how its own actions affect the world, including how they affect its own well-being. The well-being is modeled by a built-in preference function that is defined on the states of the focus. In learning to act, the system needs to store favorable action sequences in memory and to assign positive and negative preferences to previously indifferent states. In the most advanced form of learning--namely, imitation--the system uses itself to model the behavior of others of its kind. [Available as Chapter 10 of the book *Sparse Distributed Memory* and as RIACS Technical Report 88.14, "The Organization of an Autonomous Learning System."]

NEW PATTERN-COMPUTER ARCHITECTURES

Louis Jaeckel

Jaeckel has filed a patent application (No.260,256) proposing an alternative design for an SDM, with some possible hardware implementations. This design is a member of a more general family of extensions of the basic SDM architecture. Jaeckel is now preparing a continuation-in-part (i.e., ammendment) to broaden the claims of that patent application. The continuation-in-part, describing a class of alternative SDM designs, will be filed as soon as the work is finalized. Having a repertoire of design alternatives is important for improving SDM performance, for dealing with correlated data, and for use in specific applicaitons areas. This work leads naturally into more general studies of associative memory. Jaeckel has set SDM in a broader context and plans to develop a theory of sparse associative memories, both distributed and non- distributed memories. This work is ongoing and will be reported in 1989.

LARGE-SCALE SIMULATION

David Rogers

Rogers has programmed the first large-scale simulator of SDM on a Thinking Machines CM-2 with 16K processing elements. This work allows experimenting with an SDM having as many as 250,000 physical memory locations with 256-bit addresses (A one-million-location SDM could be programmed on a full-size CM-2). This work was described in two reports. [D. Rogers, "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm well-suited to the Connection Machine," RIACS Technical Report 88.32; D. Rogers, "Kanerva's Sparse Distributed Memory: An associative memory algorithm well-suited to the Connection Machine," accepted for *International Journal of High Speed Computing*.]. In 1989 Rogers will implement a new simulator of the SDM on the CM-2 using the beta release of Version 5 software. This work will be reported in 1989, a users manual will be published, and the CM-2 software will be documented and made available to the research community.

STATISTICAL SEARCH WITH SDM

David Rogers

Rogers began work on an SDM-based statistical-search technique for prediction, which incorporates concepts of John Holland on genetic classification. The objective is to predict a phenomenon of interest, given observations of state variables of the world. The fundamental idea is to exploit SDM's ability to probe an astronomically large address space with a relatively sparse, random selection of addresses. In this preliminary study, observations of state variables are encoded as addresses, and the phenomenon of interest is encoded as a binary number: "1" to indicate occurrence of the phenomenon, "0" to indicate non-occurrence. Each observation of the world thereby yields an address and a value that may be stored in a single counter at each selected location. Rogers has shown how a sparse distributed memory that has been trained on numerous such observations of the world may, in principle, yield general regions of the address space that correlate in some degree with the phenomenon of interest. Next, a genetic-classification scheme is used to search through these general regions to find specific regions of the address space with improved correlations. Addresses located in the specific regions may serve as predictors of the phenomenon. For example, previously unsuspected environmental conditions may yield a fairly reliable predictor for rain. This ongoing work has been reported in a conference paper and a technical report. [D. Rogers, "Statistical Prediction with Kanerva's Sparse Distributed Memory," in *Proceedings of 1988 Neural Information Processing Systems Conference*; D. Rogers, "Statistical Prediction with Kanerva's Sparse Distributed Memory," RIACS Technical Report 89.2.]

RECOGNITION OF 2-D SHAPES

Pentti Kanerva
Bruno Olshausen

To study the interplay of a visual preprocessor and the memory, we are testing an SDM-based technique for recognition of two-dimensional shapes. Our basic idea is to allow a shape to define a tangent field and then encode the field as a long bit-string in such a way that moderate perturbations of the shape do not alter the encoding much in terms of Hamming distance. This provides a natural representation of the shape suitable for use in a sparse distributed memory. The long bit-strings that arise from the tangent field are handled easily by the SDM. This technique has been tested on hand-drawn characters, using Olshausen's simulator running on a Sun workstation, and the initial results look promising. Early stages of the work were presented at the NASA Vision Science and Technology Workshop at Ames Research Center in December of 1988. We will continue developing the technique in 1989 by drawing upon ideas brought forth in Olshausen's survey (see below), and we plan to test it on a large database of complex shapes.

VISUAL ENCODING OF CHARACTERS

Louis Jaeckel

An autonomous system based on sparse distributed memory requires sensors to sample the state of the world. The sensory data must be encoded in ways that take advantage of the memory's capabilities. The encoding of visual images is one example of this problem. Continuing work begun in late 1987, Jaeckel developed some methods for representing and encoding a simple class of images (which includes printed characters), and two possible ways of implementing an SDM based on these encoding methods. In Jaeckel's approach, an image is segmented into *generalized arcs* that may be represented by points on a bounded five-dimensional surface, which, surprisingly, has a twist like a Moebius strip. The implementations use different strategies for mapping these sets of points on the Moebius strip onto points in the address space of a sparse distributed memory. Jaeckel also developed a small-scale simulator for one of his implementations and has been using it to study the performance of the memory. The purposes of this work are to gain insight into the problems of encoding more complex visual information, and to provide an example of encoded sensory data so that various aspects of SDM's performance can be tested and so that alternative SDM designs can be compared. Early draft reports describing the encoding methods and the SDM implementations were completed. This work will be reported in 1989.

VISUAL SHAPE REPRESENTATION

Bruno Olshausen

To use sparse distributed memory to recognize visual images, it is necessary that the images be preprocessed in order to reduce the amount of data that must be stored. The human visual system does a great deal of preprocessing upon the retinal image in order to extract the most essential information for further analysis and storage. This processing is realized by neurons interacting in a highly organized and massively parallel way. Exactly how the processing is accomplished is not yet known. Recent advances in neuroscience have revealed more of the details of brain architecture and function; these findings have been applied by many researchers in the setting of massively parallel to design artificial neural networks for pattern recognition. To determine what results of these efforts might be applicable to the problem of preprocessing images for an SDM, Olshausen has conducted an extensive survey of the research literature on biological and machine vision and neural networks. He has produced an annotated bibliography and summary of this information, and continues adding to our database of such references. [B. Olshausen, ''A Survey of Visual Preprocessing and Shape Representation Techniques,'' RIACS Technical Report 88.35, November 1988.]

LANGUAGE PROCESSING -- SPEECH

Douglas G. Danforth

A speech laboratory was proposed, investigated, and purchased. It now consists of a high-speed signal-processing board (DSP-16) made by Ariel Corporation incorporated into a SUN Microsystems 386i workstation with 8 megabytes of RAM, 91 megabytes of hard disk, and a 19-inch monitor. The lab also has a microphone and preamplifier for direct speech input as well as speakers for output. The TIMIT acoustic phonetic-speech database distributed by the US National Institute of Standards and Technology is now installed and available to the project; it will serve as a basis for research on continuous-speech transcription. Software for the lab consists of a driver, a debugger, and a Fast Fourier Transform package for the DSP-16. The language MAINSAIL was purchased for the SUN 386i to provide a flexible and convenient development environment for speech research. To date, a sparse distributed memory simulator with folds, an MSDOS-to-UNIX assembly-language translator, and a letter-sequence frequency- count routine have been written in MAINSAIL on the 386i. The frequency-count routine has been run against a 23,000-word English dictionary as a preliminary study prior to constructing a full speech-transcription system using sparse distributed memory.

TEXT TO PHONEME TRANSLATION

Umesh Joglekar

Reading aloud from printed text is among the problems that cannot be easily solved by conventional computing methods. The widely publicized NETtalk program for translating written English into phonemes (Sejnowski and Rosenberg, 1986) demonstrated that it is possible for a parallel network of computing units to be trained to form internal representations of the regularities in a training set. Because of the irregularity of English spelling, the task is far from trivial. The NETtalk experiment opens the door to a host of questions, such as: What kind of network architecture is really suited to problems of this kind? What learning strategies could be used? Is it possible to devise a system based on distributed representations that will be able to not only form abstractions of regularities in the training set but also translate these to other test data to show equally good generalization? To confirm our hypothesis that SDM may be used to solve this interesting practical problem, Umesh Joglekar began began in 1986 to develop an experimental SDM-based system for mapping text to phonemes and has simulated the system on a multiprocessor computer and on a Sun workstation. In 1988 Joglekar completed development of the system and concluded a series of experiments yielding results that are in good agreement with those reported by Sejnowski and Rosenberg in 1986. Joglekar completed a draft Masters thesis for UC Santa Barbara and a draft RIACS Technical Report describing this work, both of which will be published in 1989.

CEREBELLAR MODELING

Egon Loebner
James Keeler
David Rogers

An ongoing collaboration including Coe Miles-Schlichting of NASA Ames Information Sciences Division, Egon Loebner of Hewlett-Packard, and James Keeler of MCC was initiated to study the current human neuroscience literature to learn about motor control and the processing of sensory data in mammals. Our present working model goes beyond the view of the cerebellum as an associative memory; from an engineering standpoint the cerebellum can be viewed as a hybrid between adaptive nonlinear controller, active-memory, and processor.

Our motivation for investigating the cerebellum is twofold. First, the cerebellum is involved in planning and fine-tuning complex physical movements, as in walking, dancing, swimming, and hence we may learn something about coordination and control of autonomous systems (for example, robots). Second, the architecture of Kanerva's sparse distributed memory is similar to that of the cerebellum.

Loebner has reported an early stage of the work. [E. Loebner, "Intelligent Network Management and Functional Cerebellum Synthesis," to appear in *Proceedings of IEEE COMPCON SPRING 89*, February 1989] The efforts in 1989 will focus on building a functional and mathematical model that is consistent with existing neurophysiological data. The work will be reported as it progresses.

DIGITAL PROTOTYPE OF SDM

Michael Flynn
Philip Hickey

The digital prototype of a sparse distributed memory has been completed at Stanford and is undergoing final testing before shipment to RIACS. We began the development of system calls and diagnostics for the prototype. The report describing the prototype, which has been published by Stanford University [Technical Report CSL-TR-87-338, "Sparse Distributed Memory Prototype: Principles of Operation"]. Hickey augmented the report to include more details about the software [RIACS Technical Report TR-88.12].

APPLICATIONS OF SDM

Alvin Surkan

Surkan visited RIACS to continue work with an SDM simulator he had written at the University of Nebraska, and to develop ideas for a general-purpose connectionist network, written in APL on a Macintosh computer, for the study of multilayered feed-forward designs. This work will be reported in 1989 as research memoranda: "Some Observations from Experiments Exploring Pattern Computing for the Classification of Personal Profiles," and "Architectural Considerations in Implementing Layered Connectionist Problem Solving Systems."

PUBLICATION LIST 1988

BOOKS AND BOOK CHAPTERS

1. **Pentti Kanerva,** *Sparse Distributed Memory*, Bradford Books, MIT Press, Cambridge, MA. 1988.

REVIEWED PERIODICALS

2. **James D. Keeler**, "Comparison Between Kanerva's SDM and Hopfield-type Neural Network Models," *Cognitive Science 12* (1988), 299-329.

3. **David Rogers,** "Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited to the Connection Machine", *International Journal on High Speed Computing,* accepted.

REVIEWED CONFERENCES

4. **Pentti Kanerva,** "Computing with Very Large Patterns." "In Parallel Models of Intelligence: How Can Slow Components Think So Fast?" *Proceedings of the 1988 AAAI Spring Symposium Series,* Menlo Park, 163-169.

5. **Pentti Kanerva,** "Cerebellar-Model Associative Memory as Generalized Random-Access Memory", *Proceedings of IEEE COMPCON SPRING 89,* Session 21B, Cerebellar Models of Associative Memory, San Francisco CA, (February 1989), accepted.

6. Egon Loebner (Hewlett-Packard Laboratories), "Intelligent Network Management and Functional Cerebellum Synthesis", *IEEE COMPCOM SPRING 89,* Session 21b, Cerebellar Models of Associative Memory, San Francisco CA, (February 1989), work in collaboration with SDM group, accepted.

7. **David Rogers,** "Statistical Prediction with Kanerva's Sparse Distributed Memory," *Proceedings of the 1988 Neural Information Processing Systems Conference,* Denver, CO, (November 1988) IEEE, to appear.

OTHER PUBLISHED ARTICLES

8. M. J. Flynn, **Pentti Kanerva,** B. Ahanin, N. Bhadkamkar, P. Flaherty, and P. Hickey, "Sparse Distributed Memory Prototype: Principles of Operation", *Stanford: Computer Systems Laboratory Report No. CSL-TR-87-338,* (February 1988). 87 pp.

TECHNICAL REPORTS

1. **David Rogers,** *Using Data Tagging to Improve the Performance of Kanerva's Sparse Distributed Memory,* TR 88.1, January 1988, 32 pp.

2. **Louis Jaeckel,** *Two Alternate Proofs of Wang's Lune Formula for Sparse Distributed Memory and an Integral Approximation,* TR 88.5, February 1988, 33 pp.

3. **Pentti Kanerva,** *The Organization of an Autonomous Learning System,* TR 88.14, May 1988, 22 pp.

4. **David Rogers,** *Kanerva's Sparse Distributed Memory: An Associative Memory Algorithm Well-Suited to the Connection Machine,* TR 88.32, November 1988, 19 pp.

5. **Bruno A. Olshausen,** *A Survey of Visual Preprocessing and Shape Representation Techniques,* TR 88.35, November 1988, 59 pp.

PRESENTATIONS 1988

1. **Umesh Joglekar**, "Learning to Read Aloud," *First Annual International Neural Network Society Meeting*, Boston, (September 1988). (Poster presentation.)

2. **Pentti Kanerva**, "Parallel Models of Intelligence," AAAI Spring Symposium Series at Stanford (March 1988). (Speaker and panelist.)

3. **Pentti Kanerva**, "Computing with Very Large Patterns," NASA Ames Research Center, Machine Learning Reading Group (April 12, 1988).

4. **Pentti Kanerva**, "Understanding Information Processing in Animals as a Way to building Intelligent Robots." First Joint Technology Workshop on Neural Networks and Fuzzy Logic. Johnson Space Center, Houston, Texas (May 3-4, 1988).

5. **Pentti Kanerva**, "Adjusting to Variations in Tempo in Sequence Recognition," *First Annual International Neural Network Society Meeting*, Boston (September 1988). (Poster Presentation.)

6. **Bruno Olshausen**, "Two-Dimensional Shape Recognition Using Sparse Distributed Memory," *Vision Science and Technology Workshop*, Ames Research Center, Moffett Field, California (November 30 - December 2, 1988).

7. **Michael Raugh**, "Sparse Distributed Memory: The Project," *Vision Science and Technology Workshop*, Ames Research Center, Moffett Field, California (November 30 - December 2, 1988).

8. **David Rogers**, "Sparse Distributed Memory," *Salishan Conference*, at Salishan Lodge, Gleneden Beach, Oregon (March 1988).

9. **David Rogers**, "Prototyping Kanerva's Sparse Distributed Memory," *Scientific Applications of the Connection Machine Conference*, Ames Research Center, Moffett Field, California (September 12-14, 1988).